

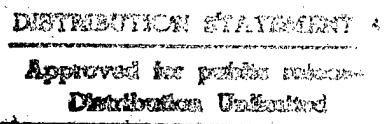
Microlanguage-Based Specialization

Technical Report for October – December 1996

Abstract

We summarize the progress made during the period July through September 1996 for the DARPA contract "Microlanguage-Based Specialization", part of the Synthetix Project. The main objective of the project is to design and implement microlanguages to support *directed* specialization of operating system kernels, to complement the *inferred* specialization in the Synthetix project. The key idea is to design a microlanguage tailored for each specific area that captures deep application semantics through a simple syntax. We are making significant progress in the construction of the Synthetix specialization toolkit and the design of the first family of microlanguages for distributed multimedia support. We also report on the progress made in the context of the entire Synthetix project.

DARPA Order Number: D007 and D308 (and B752)
Name of contractor: Oregon Graduate Institute of Science & Technology
Contract Number: F19528-95-C-0193
Principal Investigator: *Calton Pu*
Phone: (503) 690-1214
FAX: (503) 690-1553
Email: calton@cse.ogi.edu
URL: <http://www.cse.ogi.edu/~calton>
Project Title: Microlanguages (Synthetix Project)



Sponsored by
Defense Advanced Research Projects Agency, CSTO/ITO
DARPA Order No. D007 and D308
Issued by ESC/ENS under contract F19528-95-C-0193

The views and conclusion contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the Defense Advanced Research Projects Agency or the U.S. Government.

19970122 049

Contents

1 Task Objectives	1
1.1 Summary of Project Objectives	1
1.2 Synthetix Deliverables	1
1.3 Microlanguage Deliverables	2
2 Technical Problems	2
3 General Methodology	3
4 Technical Results	5
4.1 Papers and Publications	5
4.2 Presentations in This Period	7
4.3 Contract Deliverables	8
4.3.1 Oct-Dec 1995	8
4.3.2 Jan-Mar 1996	8
4.3.3 Apr-June 1996	9
4.3.4 July-September 1996	10
4.3.5 October-December 1996	11
5 Important Findings	12
6 Significant Hardware Developments	12
7 Special Comments on Industrial Collaboration	12
8 Implication for Further Research	13
A Attached Papers	14
Bibliography	14

1 Task Objectives

1.1 Summary of Project Objectives

This contract (Microlanguage, number F19528-95-C-0193) strengthens and complements the Synthetix project grant (DARPA/ONR grant N00014-94-1-0845). This research is also augmented with another DARPA grant in the Information Survivability program (Immunix, DARPA grant F30602-96-1-0331). We will report the progress and achievements of the project as a whole, distinguishing the results from each contract funding as appropriate.

For the Synthetix grant, the concrete deliverables fall into three categories: a methodology for applying it to existing operating system kernels, tools to support incremental specialization, and kernel code that demonstrates its use in practice. Similarly, for the Microlanguage contract, the concrete deliverables fall into three groups: techniques, microlanguages plus their associated software tools, and experimental systems that use microlanguages.

Synthetix specialization is primarily targeted towards situations where the systems software can *infer* the kind of specialization to apply. The goal of microlanguages is to expand the power and applicability of specialization techniques and toolkit to situations where systems programmers and application programmers may *direct* the specialization process. Some of the basic specialization techniques and tools apply equally to inferred and directed specialization, but many methods and tools need to be refined and redesigned to take into account both kinds of specialization.

1.2 Synthetix Deliverables

We will develop a methodology for the design and implementation of operating systems (and their components) that use incremental specialization. Concretely, this methodology will outline a series of steps that will guide kernel developers both in the development of new kernels from scratch, and in the modification of existing kernel code. The methodology will apply to existing monolithic kernels and micro-kernels.

We will develop a toolkit that includes a C preprocessor and compiler with support for: (a) fine-grain modularity, (b) the specification of invariants and guards, and (c) the generation of templates that can be instantiated post compile time. In addition, the toolkit will also include a run-time kernel supporting dynamic code generation, dynamic linking, and possibly some code optimization such as constant folding and loop unrolling.

Another important part of Synthetix specialization toolkit is the support for fine-grain adaptation using software feedback. The software feedback toolkit will include (1) software implementation of generic filters to be used in the kernels, (2) software tools for filter design, development, and testing, (3) composition tools to combine elementary filters into more sophisticated filters, and (4) guard programs that detect input oscillation beyond the specified filter range.

The specialization toolkit will be demonstrated in production operating system kernels. Our initial experiments were conducted on the HP-UX, the HP commercial version of Unix. Current experiments are being conducted on Linux, a public domain Unix operating system.

1.3 Microlanguage Deliverables

The first group of Microlanguage deliverables includes the expanded research methodology on specialization and microlanguage design strategies. The design of a microlanguage is complemented by microcompilers that generate code and microengines that execute the compiled microprograms. We will develop techniques to integrate the specialized underlying operating system kernel and the specialized microengine, to minimize the overhead of running the microprograms.

The second group of Microlanguage deliverables includes the microlanguages themselves and the software tools that support the microlanguages such as microcompilers and microengines. We have been using a DARPA-community compiler toolkit (SUIF of Stanford) for the development of incremental specialization tools in the Synthetix project. Our main contribution will be the techniques and tools to help microlanguage designers to handle deep application semantics.

The third group of Microlanguage deliverables consists of system components that we use to experiment with microlanguages, evaluate their performance as well as flexibility, and to demonstrate the advantages of microlanguages. For example, the experimental evaluation of first microlanguage for file systems may be conducted on a simplified version of Unix File System.

We plan to have deliverables from each group (with increasing degrees of sophistication, robustness, and variety) for the Base Period, Option 1, and Option 2 of the Microlanguages contract.

2 Technical Problems

Modern operating systems have been growing in size and complexity due to the constant pressure for additional functionality. As the variety of applications widens, and hardware platforms become increasingly powerful, the operating system has been used for different purposes. Consequently, operating system code has been stretched far beyond its original intent. For example, file systems optimized for block size access typically do not support byte-by-byte access efficiently. Although libraries such as `stdio` keep work outside the kernel, the result is the growing size and complexity of both the kernel code and the systems software related to the kernel.

To aggravate the problem, kernel operations are encapsulated within blackbox boundaries, so users cannot customize the kernel very easily. Specialization of the operating system

kernel has been developing as a promising technique in a number of projects such as Synthesis [22, 17], *x*-kernel [18], Spin [2], and Synthetix [8, 19]. However, specialization has encountered two limitations in its wide application. First, it has been applied successfully in specific domains, but each domain seems different enough to require a large new effort. Second, efforts to make customization easy, such as Spin, have yet to address the concerns of quality control, interference with other kernel modules, maintainability, and system evolution.

We propose to apply discipline to specialization through *microlanguages*, meta-interface languages that characterize how an interface is used, and tailored for each application. Each microlanguage specifies all the specialized execution paths in an appropriate systems software component, e.g., file system, network protocol, and a client/server interface for an application area. On the client’s side, a sufficiently rich microlanguage provides a high level declarative way to express customization, rather than by writing code in (third-generation) system implementation language that then has to be “sanitized” and inserted into the kernel. On the systems side, kernel designers determine the scope of specialization during the design and implementation of microlanguages.

The microlanguage approach can be explained intuitively by an analogy with the highway system. Before the highway system was invented, driving long distance was an adventure that connected from one local road to another local road. Trucking was a business that contained a large amount of uncertainty, since road conditions were unpredictably variable (e.g., due to local maintenance constraints) and the long distance traffic interfered with many local communities. Execution paths within operating systems have similar patterns to such road traffic patterns.

In our vision, instead of directly providing the user with a third-generation programming language without restrictions to customize the kernel, we introduce specialized and limited microlanguages that allow users to program specific functions in the kernel (and higher levels) in a safe, structured and disciplined way. The same way most drivers rather use existing roads (if available) instead of building their own, we think that most programmers prefer to specify what they want instead of writing new code that must be debugged and maintained. Microlanguages are a structured and disciplined way of “programming” a highly-parameterized (sub-)system. As the number of parameters and knobs increase in the system, *x*-kernel style microprotocol approaches become more difficult, since static parameters alone may not conveniently describe all the situations.

3 General Methodology

Our approach can be divided into four stages. The first stage is the design of a microlanguage for a chosen system software component (or layer). The design goal of a microlanguage is to allow a minimal set of primitives, parameters, and constraints to be expressed in a controlled

and simple way. Using our highway analogy, the designer of a microlanguage is laying out a highway system for that area. It is easy to see that we want to reduce the investment made in building the highway system, while covering the area well.

The design of microlanguages depends on the application semantics of the chosen area, and also on the current usage of the application. The same way a good highway system serves better an area more densely populated, we will support finer granularity control for execution paths more frequently traversed. Since the usage patterns may change over time, it is important to make microlanguages extensible, since it will be maintained continuously, just like highway systems. By focusing on parameters, declarations, and constraints, we want to maximize the extensibility of microlanguages, since these constructs tend to have high orthogonality.

The second stage of our research is to implement a static microcompiler and the run-time microengine (a better word might have been *microkernel*, but that term is taken). Our microengines execute microprograms, like the hardware microengines. The main technique used in this stage is *specialization*, used extensively in current operating system research, including Synthesis project at Columbia University, Synthetix project at Oregon Graduate Institute, Spin project at University of Washington, and Mach 4.0 project at University of Utah. Specialization has been shown to be very effective in the optimization of kernel calls. However, disciplined use of specialized kernel code has been a challenging problem. In our proposed research, the microcompiler generates code that directs the underlying kernel specialization, and the microengine generates the specialized code at run-time, if necessary, for best results.

The third stage is to write higher level software and evaluate the microlanguage fitness to the application area and the improvements achieved as compared to software without the microlanguage. In the past, specialization of operating system kernel calls have been evaluated based primarily on microbenchmarks, comparing kernel calls before optimization and after optimization. We have done the same in the Synthetix experiment on HP-UX [19].

However, microbenchmarks have their limitations. Specifically, they show only the effects on the particular kernel call, isolated from the rest of the system. In reality, it is the end-to-end performance that is of most interest to the user. Since microlanguages can facilitate cross-layer optimization, we will develop suitable benchmarks that will compare effectively the performance of microlanguage-based systems and those without microlanguages.

The fourth stage is to use the evaluation results to refine the microlanguage and its attendant tools such as microcompiler and microengine. Having done the experimental evaluation of systems with microlanguages, we will use the results to refine the experiment and the system. In particular, we will extend/reduce the microlanguage appropriately, according to the evaluation. Using our tools, we will then modify the microcompiler and microengine to re-evaluate our system to verify the impact of our refinements.

4 Technical Results

During the period covered by this report, we have made significant advances in our project. We first outline a selection of our publications and relevant papers written during this period.

4.1 Papers and Publications

During the Oct–Dec 1995 period:

1. Jonathan Walpole presented our paper [19] at the 1995 ACM Symposium on Operating Systems Principles (SOSP). The paper described our experience with applying specialization techniques systematically to a commercial operating system (HP-UX). In this paper, we have demonstrated the potential gains of specialization in performance and modularity on the one hand, and on the other hand, the difficulties with verifying the correctness of the specialized code, requiring powerful software tools for the large scale application of specialization. The SOSP is the bi-annual premier conference in operating systems research.
2. Shanwei Cen presented a demonstration [4] of our real-time distributed MPEG video/audio player at the 1995 ACM International Conference on Multimedia. The demo showcased our software-feedback based software, which has been freely available on the Web since April 1995.
3. Our paper [12] on “Adaptive Methods for Distributed Video Presentation” has appeared in a special Symposium on Multimedia of the ACM Computing Surveys journal. The paper analyzes the problems and solutions for delivering real-time multimedia presentations across the Internet, with special attention to our approach based on software feedback and our demonstration software.

During the Jan–Mar 1996 period:

1. Our paper [7] on “A General Approach for Run-Time Specialization and its Application to C” has been presented in the 1996 ACM Symposium on Principles on Programming Languages (POPL). This paper describes the theoretical foundation of the C on-line partial evaluator Consel is developing in France as part of the Synthetix specialization toolkit. The POPL is the annual premier conference on programming languages research.
2. Our paper [27] on “Safe Operating System Specialization: the RPC Case Study” has been presented at the First Workshop on Compiler Support for Systems Software (WCSSS), Arizona, February 1996. This paper describes the experiments in applying specialization to the RPC code of the Chorus microkernel.

3. Our paper [6] on “A Uniform Approach to Compile-Time and Run-Time Specialization” has been presented at the 1996 Dagstuhl Workshop on Partial Evaluation, Germany, February 1996. This paper describes the general approach of the Rennes group to specialization at both compile-time and run-time.

During the Apr–June 1996 period:

1. The paper [13] on “A Wait-free Algorithm for Optimistic Programming: HOPE Realized”, by Crispin Cowan and Hanan L. Lutfiyya, was presented at the 1996 International Conference on Distributed Computing Systems, Hong Kong, May 1996.
2. Our paper [26] on “A Uniform Automatic Approach to Copy Elimination in System Extensions via Program Specialization”, by Volanschi, E.-N., Muller, G., Consel, C., Hornof, L., Noye, J. and Pu, C., was published as IRISA Research Report No. 1021, Rennes, France, June 1996. It was submitted to the 1996 Symposium on Operating System Design and Implementation, Seattle, October 1996.
3. Our paper [9] “Automated Guarding Tools for Adaptive Operating Systems”, by Crispin Cowan, Andrew Black, Charles Krasic, Calton Pu, and Jonathan Walpole was submitted to the 1996 Symposium on Operating System Design and Implementation, Seattle, October 1996.

During the July–September 1996 period:

1. Our paper [25] “Managing Adaptive Presentation Executions in Distributed Multimedia Database Systems” by H. Thimm, W. Klas, J. Walpole, C. Pu, and C. Cowan, presented in the 1996 IEEE International Workshop on Multimedia Database Management Systems, Blue Mountain Lake, New York, August 1996.
2. Our paper [3], “System Support for Mobility”, by A. Black and J. Inouye, was presented at the ACM SIGOPS European Workshop in Ireland, September 1996.
3. Our paper [11] “Specialization Classes: An Object Framework for Specialization”, by Crispin Cowan, Andrew Black, Charles Krasic, Calton Pu, and Jonathan Walpole of OGI, and Charles Consel and Eugen-Nicolae Volanschi of University of Rennes / IRISA, accepted for presentation in the 1996 International Workshop on Object Orientation in Operating Systems, Seattle, October 1996.
4. Our paper [20] “A Specialization Toolkit to Increase the Diversity in Operating Systems” by C. Pu, C. Cowan, A. Black and J. Walpole of OGI, and C. Consel of IRISA/University of Rennes, submitted for presentation in the ICMAS Workshop on Immunity-Based Systems, Kyoto, Japan, December 1996.

5. Our paper [21] "Microlanguages for Operating System Specialization" in Operating Systems" by C. Pu, C. Cowan, A. Black and J. Walpole of OGI, and C. Consel of IRISA/University of Rennes, submitted for presentation in the POPL Workshop on Domain-Specific Languages, Paris, January 1997.

During the October–December 1996 period:

1. Crispin Cowan presented the specialization class paper [11] at the International Workshop on Object-Oriented in Operating Systems, Seattle, October, 1996.
2. Veronica Baiceanu [1] and Dan Revel [23] presented their work on CPU and disk scheduling for multimedia applications at the Workshop on Resource Allocation Problems in Multimedia Systems, Washington, D.C., December 1996.
3. Calton Pu presented our ideas for improving operating system survivability [20] at the Immunity-Based Sysetms Workshop, Japan, December 1996.

Pending journal submissions:

1. We have submitted for journal publication an extended version of a recent conference paper [14]. The paper describes the formal semantics for expressing optimism, which is a general approach for parallel and distributed computation compatible with specialization.
2. We have been invited to submit our paper "Customizable Operating Systems" as a book chapter, being edited as a result of a special issue of CACM on advances on operating systems.

4.2 Presentations in This Period

1. Jonathan Walpole gave the talk on Quality of Service Specification and Adaptive Resource Management for Large Distributed Systems at the Quorum PI meeting (Dallas, November 1996).
2. Crispin Cowan visited CMU, Prof. Daniel Jackson, on December 9th, 1996, to discuss collaboration. We are actively studying the use of *lackwit*, their software analysis tool, for the static verification of guarded code. The OGI group has sent real code to be analyzed by the CMU group as the preliminary evaluation of the tool robustness for production code.

4.3 Contract Deliverables

4.3.1 Oct–Dec 1995

Specialization Methodology We have been refining the specialization methodology in our work. A preliminary version of the methodology has been described in our SOSP paper [19]. The refined methodology, which includes the use of and support for microlanguages, is under active development. This is both Synthetix and Microlanguage deliverable.

Specialization Toolkit We have been designing and implementing the specialization toolkit during this period. Our main accomplishment is the creation and refinement of the concept of *specialization class*. Each concrete specialization is an instance of a general type of specialization, defined by the specialization class. Many tools in the specialization toolkit are based on this concept. The explicit (and machine readable) definition of specialization classes allows the software tools to communicate with each other easily. We are writing several papers on specialization classes and developing code using them. Preliminary versions of specific tools have been under test by the members of our group, including the Tempo-C on-line partial evaluator in France and the dynamic kernel module replugger at OGI. The toolkit includes deliverables from both the Synthetix grant and the Microlanguages contract.

Software Feedback We have been designing and implementing a software feedback toolkit for fine-grain adaptive distributed computation. There are two recent developments in this effort. First, conceptually, we are modeling software feedback as a fine-grain specialization technique. This way, we will be able to use the specialization tools to support software feedback, and integrate the software feedback toolkit into the specialization toolkit easily. Second, concretely, we are re-implementing the distributed multimedia video/audio player to use the software feedback toolkit. This is a Synthetix deliverable.

Microlanguage Development We have started the development of our first microlanguage for file systems. From the conceptual side, this effort is based on recent work on the specialization of file systems such as our SOSP paper [19] and other research papers on choosing the appropriate caching strategy (e.g., [16]) for an application. From the systems building side, this work is integrated with the construction of the specialization toolkit. We are re-designing the specialization tools to make them extensible and add hooks supporting microlanguages from the beginning. This is a Microlanguages deliverable.

4.3.2 Jan–Mar 1996

Specialization Methodology We have been refining the specialization methodology in our work. In this quarter, we have been designing the next generation specialization experiments in the context of Quality of Service support for multimedia and dynamic reconfiguration of mobile networks. This is both a Synthetix and Microlanguage deliverable.

Specialization Toolkit We have been refining the concept of specialization class. Each concrete specialization is an instance of a general type of specialization, defined by the spe-

cialization class. Currently, the specialization class is considered the “assembly language” of specialization, into which the microprograms will be translated. Components completed for alpha testing include the Replugger [9] the type-based guard checker, and the preliminary version of the Tempo-C compiler [7]. The memory run-time guard checker is under development at OGI. The specialization toolkit will include deliverables from both the Synthetix grant and the Microlanguages contract.

Software Feedback We have redesigned the software feedback toolkit. Currently, it consists of a component library being written in C++, composition tools, and guarding tools built using the specialization toolkit. The component library will be illustrated using a new demonstration program that has enhanced functionality compared to the distributed multimedia MPEG player we released in May 1995. This is a Synthetix deliverable.

Microlanguage Development We have started the development of our first family of microlanguages. The family is an integrated set of four microlanguages to describe and program multimedia applications. At the top level, it will be a Quality of Service (QoS) specification microlanguage, based on Richard Staehli’s dissertation work [24] under Jonathan Walpole. The QoS specification is then translated into a media-aware scheduling microlanguage, which knows about the compression method, time binding, frame dependencies, and other resource management issues. The data access part of the scheduling microlanguage is then translated into a low level file system microlanguage, which knows about the details of file representation and access. The file system microlanguage is finally translated into the specialization classes being written for the kernel experiments, including read-ahead, buffer cache usage and management, and I/O system performance. This is a Microlanguages deliverable.

4.3.3 Apr–June 1996

Specialization Methodology We have been refining the specialization methodology in our work. Using the plans for experimentation, we have been designing and implementing the specialization and software feedback toolkit as described below. A new experiment on automated elimination of copying has been reported [26]. This is both a Synthetix and Microlanguage deliverable.

Specialization Toolkit We have been refining the concept of specialization class. In this quarter, we have completed two new papers. The first paper [9] describes the first part of the toolkit. The second paper [11] describes the current refinements of specialization classes. Besides the Replugger [9], the type-based guard checker, and the preliminary version of the Tempo-C compiler [7], we have continued the development of the memory run-time guard checker. The specialization toolkit will include deliverables from both the Synthetix grant and the Microlanguages contract.

Software Feedback We have been implementing the redesigned the software feedback toolkit. Currently, it consists of a component library being written in C++, composition

tools, and guarding tools built using the specialization toolkit. The component library will be illustrated using a new demonstration program that has enhanced functionality compared to the distributed multimedia MPEG player we released in May 1995. This is a Synthetix deliverable.

Microlanguage Development We have continued the development of our first family of microlanguages. The design has been refined using our experience with the implementations of the multimedia player and the file system specialization experiment. The family is an integrated set of four microlanguages to describe and program multimedia applications. At the top level, it will be a Quality of Service (QoS) specification microlanguage. The QoS specification is then translated into a media-aware scheduling microlanguage. The data access part of the scheduling microlanguage is then translated into a low level file system microlanguage. The file system microlanguage is finally translated into specialization classes. Mr. Thimm is working on the high-level microlanguage for Quality of Service specification and software integration of Synthetix player and GMD/IPSI's multimedia storage support programs. This is a Microlanguages deliverable.

4.3.4 July–September 1996

Specialization Methodology We have been refining the specialization methodology in our work. Using the plans for experimentation, we have been designing and implementing the specialization and software feedback toolkit as described below. This is both a Synthetix and Microlanguage deliverable.

Specialization Toolkit We have been refining the concept of specialization class. In this quarter, we have completed two new papers. The introductory paper [11] describes the current refinements of specialization classes. In addition to the Replugger [9] and the Tempo-C specializer [7], we have continued the development of the type-based guard checker and the memory run-time guard checker. The specialization toolkit will include deliverables from both the Synthetix grant and the Microlanguages contract.

Software Feedback We have been implementing the redesigned the software feedback toolkit. Currently, it consists of a component library being written in C++, composition tools, and guarding tools built using the specialization toolkit. The component library will be illustrated using a new demonstration program that has enhanced functionality compared to the distributed multimedia MPEG player we released in May 1995. This is a Synthetix deliverable.

Microlanguage Development We have continued the development of our first family of microlanguages. The design has been refined using our experience with the implementations of the multimedia player and the file system specialization experiment. The family is an integrated set of four microlanguages to describe and program multimedia applications. At the top level, it will be a Quality of Service (QoS) specification microlanguage. The QoS specification is then translated into a media-aware scheduling microlanguage. The data

access part of the scheduling microlanguage is then translated into a low level file system microlanguage. The file system microlanguage [21] is finally translated into specialization classes. Mr. Thimm is working on the high-level microlanguage for Quality of Service specification and software integration of Synthetix player and GMD/IPSI's multimedia storage support programs [25]. This is a Microlanguages deliverable.

4.3.5 October–December 1996

Specialization Methodology We have been refining the specialization methodology in our work. Using the plans for experimentation, we have been designing and implementing the specialization and software feedback toolkit as described below. This is both a Synthetix and Microlanguage deliverable.

Specialization Toolkit The paper on specialization classes [11] was presented in this quarter. In addition to the Replugger [9] and the Tempo-C specializer [7], we have continued the development of the type-based guard checker and the memory run-time guard checker. All tools are being tested and used in our experimental work. The specialization toolkit will include deliverables from both the Synthetix grant and the Microlanguages contract.

Software Feedback We have been implementing the redesigned the software feedback toolkit. Currently, it consists of a component library being written in C++, composition tools, and guarding tools built using the specialization toolkit. The component library will be illustrated using a new demonstration program that has enhanced functionality compared to the distributed multimedia MPEG player we released in May 1995. This is a Synthetix deliverable. A Master's thesis was completed [15], on the "Design of a Multimedia Player with Advanced QoS Control". This is a Synthetix deliverable.

Microlanguage Development We have continued the development of microlanguages. The design has been refined using our experience with the implementations of the multimedia player and the file system specialization experiment. The file system microlanguage [21] is under continued development. We are also exploring a variety of scheduling approaches [1, 23] and developing a CPU scheduling microlanguage to be combined with file prefetching, to help achieve multimedia QoS requirements. We are also designing a microlanguage for mobile applications.

Experiments Using Microlanguages We have started experimenting with microlanguages for the file system prefetching and will be experimenting with the other scheduling microlanguages when they are implemented in the prototype form. In the mobile area, we are focusing on cost management of wireless connections. This is a Microlanguage deliverable.

5 Important Findings

The Microlanguages contract consists of a combination of efforts from the operating systems side and from the programming languages side. We have summarized our important findings in two papers, both appearing in the most important conference of the area. In our SOSP paper [19], our most important findings in operating systems are:

1. Systematic specialization is feasible and worthwhile (significant performance gains in the kernel calls studied). Our current work on specialization classes [11] follows this thread.
2. Large-scale specialization requires sophisticated and easy-to-use software tools. See the discussion below.

In our POPL paper [7], our most important findings in the programming languages area are:

1. It's possible to build an on-line partial evaluator (Tempo-C) to support specialization.
2. It's possible to apply the partial evaluator to significant pieces of operating system code written in (cleaned-up) C. Our current experiments [26] follow this thread.

Our experimental work [6, 27, 12, 4] has validated our research ideas and provided directions to its continued development.

Our toolkit construction work is making good progress. The specialization toolkit consists of alpha test of the Replugger [9], Tempo-C compiler [7], and type-based guard checker [10], with the memory run-time guard checker being integrated into the Linux kernel. The software feedback toolkit consists of a alpha test of the component library and a re-implementation of the distributed multimedia MPEG player [5].

6 Significant Hardware Developments

The Synthetix project is a software development project. We have not planned any hardware developments. During this period of contract performance, no unexpected hardware developments were necessary.

7 Special Comments on Industrial Collaboration

In the Synthetix project, we have maintained an active collaboration with industry. Currently, we have the following active collaborations:

- Tektronix: distributed multimedia application support based on software feedback. We are receiving funding from Tektronix for this collaboration.

- Xerox PARC: use of reflection in specialization and the application of specialization as reflection.
- Intel: application of specialization to distributed multimedia and architectural support for specialization. We received funding in 1996 for this collaboration.

8 Implication for Further Research

The project is ongoing and we are working on the following aspects of the research:

- Specialization Methodology: We are continuing to refine the methodology as we develop the specialization and software feedback toolkits. We are redoing the specialization experiments using the toolkit, as we use the experience to refine the toolkits. Ongoing experiments include specialization of Unix `vmalloc` for performance, and specialization of signal handlers for modularity, and random specialization for generic operating system modules for survivability.
- Specialization Toolkit: We are using toolkit component in our experimental work. Preliminary versions of specific tools continue to undergo testing by the members of our group. Both the TypeGuard and the MemGuard are being tested at OGI. We are using Tempo in France and OGI for performance and survivability. We are also evaluating several software tools developed elsewhere, e.g., the lackwit of CMU, developed under the EDCS program.
- Software Feedback: The new software feedback toolkit continues to be developed, including the component library, composition and guarding tools, and the demonstration distributed MPEG multimedia player program. The composition and guarding tools will use the specialization toolkit wherever feasible. This under active development.
- Microlanguage Development: We are developing our first family of microlanguages, which consists of a QoS specification microlanguage, a resource scheduling microlanguage, a file system microlanguage, and appropriate specialization classes in the kernel. An example of microlanguage has been designed for the file system [DSL].
- Information Survivability: We will carry out the experimental evaluation of specialization as a defense against system attack methods we have collected. Most of the attacks use the overflow of buffer allocated on the program stack. We plan to vary the allocation of buffer, both in size and location, and verify the resiliency to the known attacks.

A Attached Papers

Four papers are attached to this technical report.

1. Paper in 1996 IWOOOS [11], labeled [IWOOOS].
2. Paper in 1996 ICMAS Workshop on Immunity-Based Systems [20], labeled [IBS].
3. Two papers in 1996 Workshop on Resource Allocation Problems in Multimedia [1, 23], labeled [RAMM1] and [RAMM2].

References

- [1] V. Baiceanu, C. Cowan, D. McNamee, C. Pu, and J. Walpole. Multimedia applications require adaptive CPU scheduling. In *Proceedings of the 1996 Workshop on Resource Allocation Problems in Multimedia Systems*, Washington, D.C., December 1996.
- [2] B.N. Bershad, C. Chambers, S. Eggers, C. Maeda, D. McNamee, P. Pardyak, S. Savage, and E.G. Sirer. SPIN - An Extensible Microkernel for Application-specific Operating System Services. In *SIGOPS 1994 European Workshop*, February 1994. UW Technical Report 94-03-03.
- [3] A. Black and J. Inouye. System support for mobility. In *Proceedings of the Seventh ACM SIGOPS European Workshop*, pages 129–132. ACM, September 1996.
- [4] S. Cen, C. Pu, R. Staehli, C. Cowan, and J. Walpole. Demonstrating the effect of software feedback on a distributed real-time MPEG video audio player. In *Proceedings of the 1995 ACM International Conference on Multimedia*, San Francisco, November 1995.
- [5] S. Cen, C. Pu, R. Staehli, C. Cowan, and J. Walpole. A distributed real-time MPEG video audio player. In Tom Little, editor, *Proceedings of the 1995 International Workshop on Network and Operating System Support for Digital Audio and Video*, volume 1018 of *Lecture Notes in Computer Science*. Springer-Verlag, 1995. Also appeared in the Proceedings of the 1995 International Workshop on Network and Operating System Support for Digital Audio and Video, April 1995, New Hampshire.
- [6] C. Consel, L. Hornof, F. Noel, and E.N. Volanschi. A uniform approach to compile-time and run-time specialization. In *Proceedings of the 1996 Dagstuhl Workshop on Partial Evaluation*, Germany, February 1996.
- [7] C. Consel and F. Noel. A general approach for run-time specialization and its application to c. In *Proceedings of the 1996 ACM Symposium on Principles of Programming Languages*, Florida, January 1996.
- [8] C. Consel, C. Pu, and J. Walpole. Incremental partial evaluation: The key to high performance, modularity and portability in operating systems. In *Proceedings of the 1993 ACM Symposium on Partial Evaluation and Semantics-Based Program Manipulation*, Copenhagen, Denmark, June 1993.

- [9] C. Cowan, T. Autrey, C. Krasic, C. Pu, and J. Walpole. Fast concurrent dynamic linking for an adaptive operating system. In *Proceedings of the International Conference on Configurable Distributed Systems*, Maryland, May 1996.
- [10] C. Cowan, A. Black, C. Krasic, C. Pu, and J. Walpole. Automated guarding tools for adaptive operating systems. Technical Report OGI-CSE-95-0XX, Department of Computer Science and Engineering, Oregon Graduate Institute, May 1996.
- [11] C. Cowan, A. Black, C. Krasic, C. Pu, J. Walpole, C. Consel, and N. Volanschi. Specialization classes: An object framework for specialization. In *Proceedings of the 1996 International Workshop on Object-Orientation in Operating Systems*, Seattle, October 1996.
- [12] C. Cowan, S. Cen, J. Walpole, and C. Pu. Adaptive methods for distributed video presentation. *ACM Computing Surveys*, 27, December 1995.
- [13] C. Cowan and H. Lutfiyya. A wait-free algorithm for optimistic programming: Hope realized. In *Proceedings of the 16th International Conference on Distributed Computing Systems*, Hong Kong, May 1996.
- [14] Crispin Cowan and Hanan Lutfiyya. Formal Semantics for Expressing Optimism: The Meaning of HOPE. In *1995 Symposium on the Principles of Distributed Computing (PODC)*, pages 164–173, Ottawa, Ontario, August 1995.
- [15] Rainer Koster. Design of a multimedia player with advanced QoS control. Master's thesis, Oregon Graduate Institute of Science and Technology, Portland, Oregon, 1996.
- [16] D. Kotz. Disk-oriented I/O for MIMD multiprocessors. In *Proceedings of the First Symposium on Operating Systems Design and Implementation*, pages 61–74, Monterrey, CA, November 1994. Usenix.
- [17] H. Massalin and C. Pu. Threads and input/output in the Synthesis kernel. In *Proceedings of the Twelfth ACM Symposium on Operating System Principles*, pages 191–201, Arizona, December 1989.
- [18] S. O'Malley and L. Peterson. A dynamic network architecture. *ACM Transactions on Computer Systems*, 10(2):110–143, May 1992.
- [19] C. Pu, T. Autrey, A. Black, C. Consel, C. Cowan, J. Inouye, L. Kethana, J. Walpole, and K. Zhang. Optimistic incremental specialization: Streamlining a commercial operating system. In *Proceedings of the Fifteenth Symposium on Operating Systems Principles*, Colorado, December 1995.
- [20] C. Pu, A. Black, C. Cowan, and J. Walpole. A specialization toolkit to increase the diversity of operating systems. In *Proceedings of the 1996 ICMAIS Workshop on Immunity-Based Systems*, Nara, Japan, December 1996.
- [21] C. Pu, A. Black, C. Cowan, and J. Walpole. Microlanguages for operating systems specialization. In *Proceedings of the 1997 POPL Workshop on Domain-Specific Languages*, Paris, January 1997.
- [22] C. Pu, H. Massalin, and J. Ioannidis. The Synthesis kernel. *Computing Systems*, 1(1):11–32, Winter 1988.

- [23] D. Revel, C. Cowan, D. McNamee, C. Pu, and J. Walpole. An architecture for flexible multi-media prefetching. In *Proceedings of the 1996 Workshop on Resource Allocation Problems in Multimedia Systems*, Washington, D.C., December 1996.
- [24] R. Staehli. *Quality of Service Specifications for Resource Management of Multimedia Systems*. PhD thesis, Department of Computer Science and Engineering, Oregon Graduate Institute of Science & Technology, 1995.
- [25] H. Thimm, W. Klas, J. Walpole, and C. Pu, C.and Cowan. Managing adaptive presentation executions in distributed multimedia database systems. In *Proceedings of the IEEE International Workshop on Multimedia Database Management Systems*, pages 152–159, Blue Mountain Lake, NY, August 1996.
- [26] E.-N. Volanschi, G. Muller, C. Consel, L. Hornof, J. Noyé, and C. Pu. A uniform automatic approach to copy elimination in system extensions via program specialization. Research Report 1021, Irisa, Rennes, France, June 1996.
- [27] E.N. Volanschi, G. Muller, and C. Consel. Safe operating system specialization: the rpc case study. In *Proceedings of the First Workshop on Compiler Support for Systems Software*, Arizona, February 1996.